

# Numerical Computation of $p$ -values with *myFitter*

M. WIEBUSCH\*

*Institute for Theoretical Particle Physics,  
Karlsruhe Institute of Technology (KIT), D-76128 Karlsruhe, Germany*

## Abstract

The computation of  $p$ -values in goodness-of-fit or likelihood ratio tests is often done with simple analytic formulae, even if the validity of these formulae is far from obvious. In particular, no simple formula exists for the comparison of models that are not *nested*, in the sense that one model can be obtained from the other by fixing some of its parameters. In this paper I present a strategy for efficient *numerical* computations of  $p$ -values, which works for both nested and non-nested models and does not rely on additional approximations. These ideas have been implemented in a publicly available C++ framework for maximum likelihood fits called *myFitter* and have recently been applied in a global analysis of the Standard Model with a fourth generation of fermions.

---

\*email: wiebusch@particle.uni-karlsruhe.de

# 1 Introduction

Even though the LHC experiments have, so far, not found any clear signs for physics beyond the Standard Model (SM) they already put strong constraints on the favourite SM extensions of many theorists. The SM with a (perturbative) sequential fourth generation of fermions (SM4) is under a lot of pressure from direct Higgs searches and may be the first to be excluded. Other models with additional fermions or even some constrained versions of Supersymmetry may follow soon.

In this situation some thoughts should be spent on the methods and criteria by which we decide if a certain model is ruled out. Two well-established techniques in (frequentist) statistical analyses are *goodness-of-fit* and *likelihood ratio tests*. (For an introduction see e.g. [1] or the statistics chapter of [2].) In the former case, the minimal value of the  $\chi^2$  function is used as a test statistic, while in the latter case one uses the difference  $\Delta\chi^2$  of the minimum  $\chi^2$ -values of two models as a test statistic. In many cases, the probability density function (PDF) of the test statistic is, to a good approximation, given by the well-known  $\chi^2$ -distribution [3]. In this case the relation between the minimal  $\chi^2$  value of a model or the  $\Delta\chi^2$  value of two models and the statistical significance (*p*-value) of the corresponding hypothesis tests is described by the normalised lower incomplete gamma function.

There are, however, also many realistic scenarios where the PDF of the test statistic is not described by a  $\chi^2$ -distribution. One example is the case of likelihood ratio tests where the two models to be compared are not *nested*, meaning that one model can not be obtained from the other by fixing some of its parameters. This problem was encountered in a recent analysis of the Standard Model (SM) with a fourth generation of fermions [4]. In this analysis it is impossible to regard the SM with three fermion generations as a limiting case of the SM with four generations due to non-decoupling contributions of chiral fermions in electroweak precision observables and Higgs production and decay rates. Another case where analytical formulae for *p*-values are not reliable is the situation where some of the parameters of a model are bounded, in the sense that they are only allowed to float within a certain range. Most notably, this applies to analyses where systematic errors are treated within the *RFit* scheme [5], i.e. by introducing so-called nuisance parameters with a limited range.

When analytic formulae fail one has to resort to numerical methods, and the computation of *p*-values is no exception. The basic strategy is to generate a large sample of random *toy measurements*, whose distribution is given by the experimental errors of the observables and which is centred on the model's prediction. For each toy measurement the value of the test statistic is computed and compared to the value obtained from the actual data. With a large enough sample we can then estimate the probability that the value of the test statistic is more extreme than the observed value, and this probability is called *statistical significance* or *p-value*. Unfortunately, the computational cost of the required numerical simulations can be rather high, especially when the *p*-value is

small. In this paper I discuss some methods for improving the efficiency of numerical computations of  $p$ -values. These methods have been applied in [6], where, based on the constraints from Higgs searches and electroweak precision observables, likelihood ratio tests comparing the SM with three and four fermion generations were performed. The methods are also implemented in a publicly available code called *myFitter*, which I present in this paper.

The paper is organised as follows: in Sec. 2 I describe the general mathematical setup and the definitions of the different types of  $p$ -values. In Sec. 3 I go through the derivation of the analytical formulae for  $p$ -values and discuss their range of applicability. In Sec. 4 I explain the strategy for improving the efficiency of numerical computations of  $p$ -values. The *myFitter* code and the implementation of the methods from Sec. 4 is presented in Sec. 5. I conclude in Sec. 6.

## 2 General Setup

To fix some notations let me first introduce the general mathematical setup for goodness-of-fit and likelihood ratio tests. The discussion essentially follows the statistics review in [2], but puts a stronger emphasis on the geometric interpretation of the log-likelihood function. This geometric picture will be needed in the next sections.

Let  $\mathbf{Q} = (Q_i)_{i=1,\dots,n}$  be a vector of observables whose experimentally measured values may lie in a *sample space*  $S \subset \mathbb{R}^n$ . Assume that the true values of these observables are described by a vector  $\hat{\mathbf{Q}} \in S$ . The results  $Q_i$  of the measurements may then be regarded as  $n$  (generally not independent) random variables distributed according to a (joint) probability density function (PDF)  $\pi_{\hat{\mathbf{Q}}}(\mathbf{Q})$ . The *log-likelihood function*  $L(\hat{\mathbf{Q}}, \mathbf{Q})$  is defined as

$$L(\hat{\mathbf{Q}}, \mathbf{Q}) = -2 \ln \pi_{\hat{\mathbf{Q}}}(\mathbf{Q}) \quad . \quad (1)$$

In this setup,  $L$  is required to have the following properties:

1. For any given  $\hat{\mathbf{Q}} \in S$ , the log-likelihood function  $L$  must satisfy the normalisation condition

$$\int_S d^n \mathbf{Q} e^{-\frac{1}{2} L(\hat{\mathbf{Q}}, \mathbf{Q})} = 1 \quad . \quad (2)$$

2. For any given  $\mathbf{Q} \in S$ , and considered as a function of  $\hat{\mathbf{Q}}$ , the log-likelihood function  $L(\hat{\mathbf{Q}}, \mathbf{Q})$  must be bounded from below and have its unique absolute minimum at  $\hat{\mathbf{Q}} = \mathbf{Q}$ .
3. For any given  $\hat{\mathbf{Q}} \in S$ , and considered as a function of  $\mathbf{Q}$ , the log-likelihood function  $L(\hat{\mathbf{Q}}, \mathbf{Q})$  must be bounded from below and have its unique absolute minimum at  $\hat{\mathbf{Q}} = \mathbf{Q}$ .

The first property is simply the normalisation condition for the PDF  $\pi_{\hat{\mathbf{Q}}}(\mathbf{Q})$ . The second property allows us to regard  $L(\boldsymbol{\Theta}, \mathbf{Q})$  as a *measure of disagreement* between a theoretical prediction  $\boldsymbol{\Theta} \in S$  and a measured set of observables  $\mathbf{Q} \in S$ . This property justifies the method of maximum likelihood fits, where theory parameters are determined by minimising  $L$ .

The third property reflects the fact that the measurements of the  $Q_i$  really measure the quantities  $\hat{Q}_i$ . Without any modifications to the model, it does *not* hold in the presence of systematic errors, since a systematic error is an offset between the true value of an observable and its most likely measured value. This offset is the same each time the measurement is performed and does therefore not average out when the measurement is repeated many times. This results in a difference between  $\hat{Q}_i$  and the maximum of the distribution of the random variable  $Q_i$ . The central idea of the RFit method [5] is that systematic errors should not be treated as errors at all, but as unknown theory parameters, so-called *nuisance parameters*, that may vary within a certain range. In a way, the presence of a systematic error means that theorists and experimentalists are simply not talking about the same quantity. Since the difference between the two quantities can neither be modeled nor measured it has to be treated as an additional model parameter, but with a limited range of possible values. Thus, the third assumption *does* hold if systematic errors are treated within the RFit scheme, i.e. by introducing a nuisance parameter for each source of systematic errors.

In this generic setting, a *theory* is simply a function  $\boldsymbol{\Theta}$  that maps parameters  $\mathbf{x} = (x_i)_{i=1,\dots,k}$  from some parameter space  $X \subset \mathbb{R}^k$  into  $S$ . In the absence of systematic errors  $k$  should be smaller than the number  $n$  of observables if one wants to constrain the parameter space  $X$ . However, if all observables are affected by systematic errors, and an independent nuisance parameter is introduced in each case, the number  $k$  of parameters may indeed be larger than  $n$ . Of course, the ranges of the nuisance parameters will be limited. In any case we define the *theory manifold*  $M$  as

$$M = \{\boldsymbol{\Theta}(\mathbf{x}) | \mathbf{x} \in X\} \subset S \quad . \quad (3)$$

To quantify the disagreement between the theory and some vector  $\mathbf{Q} \in S$  of measured observables we define the *squared distance function*  $d^2$  by

$$d^2(\mathbf{Q}) = \min_{\mathbf{Q}' \in M} [L(\mathbf{Q}', \mathbf{Q}) - L(\mathbf{Q}, \mathbf{Q})] = \min_{\mathbf{x} \in X} [L(\boldsymbol{\Theta}(\mathbf{x}), \mathbf{Q}) - L(\mathbf{Q}, \mathbf{Q})] \in \mathbb{R} \quad . \quad (4)$$

The definition of  $d^2(\mathbf{Q})$  is the usual definition of the squared distance between the manifold  $M$  and the point  $\mathbf{Q}$  if the squared distance between two points in  $S$  is measured by  $L$ . Note that I will use the term “squared distance” for the values  $L(\hat{\mathbf{Q}}, \mathbf{Q})$  even though, in this general setting,  $\sqrt{L}$  is not guaranteed to have the usual properties of a distance measure, such as positivity or the triangle inequality. However, these properties are not needed for our discussion. All we require for our definitions to make sense are the properties 1 to 3 from the start of this section. In the case of Gaussian errors we

will see that  $L(\hat{\mathbf{Q}}, \mathbf{Q})$  is (up to a constant) indeed just the squared euclidean distance between  $\hat{\mathbf{Q}}$  and  $\mathbf{Q}$ .

Except for very contrived and pathological cases, there will only be one element of  $M$  with minimal distance to  $\mathbf{Q}$ . We denote these *best-fit observables* as  $\hat{\mathbf{Q}}(\mathbf{Q}) \in S$ , so that

$$L(\hat{\mathbf{Q}}(\mathbf{Q}), \mathbf{Q}) = d^2(\mathbf{Q}) \quad . \quad (5)$$

The domain  $\hat{X}(\mathbf{Q}) \subset X$  of best-fit parameters is then

$$\hat{X}(\mathbf{Q}) = \Theta^{-1}(\hat{\mathbf{Q}}(\mathbf{Q})) = \{\mathbf{x} \in X | \Theta(\mathbf{x}) = \hat{\mathbf{Q}}(\mathbf{Q})\} \quad . \quad (6)$$

In the absence of systematic errors the set  $\hat{X}(\mathbf{Q})$  will usually contain only one element, namely, the best-fit parameter vector  $\hat{\mathbf{x}}(\mathbf{Q})$ . However, if many nuisance parameters are introduced because of systematic errors there may be ambiguities. The definitions below rely on  $\hat{\mathbf{Q}}(\mathbf{Q})$  being unique, but not on  $\hat{\mathbf{x}}(\mathbf{Q})$  being unique. Note that, in order to compute  $d^2(\mathbf{Q})$  and  $\hat{\mathbf{x}}(\mathbf{Q})$  one has to go through the usual procedure of minimising the chi-square function, which, for a given vector of measured observables  $\mathbf{Q}$ , is defined as

$$\chi^2(\mathbf{x}) = L(\Theta(\mathbf{x}), \mathbf{Q}) - L(\mathbf{Q}, \mathbf{Q}) \quad . \quad (7)$$

The best-fit parameters  $\hat{\mathbf{x}}(\mathbf{Q})$  are then just the parameters that minimise the  $\chi^2$  function and  $d^2(\mathbf{Q})$  is the minimal  $\chi^2$ -value.

With these preparations we are now ready to write down the formula for the  $p$ -value  $p(\mathbf{Q})$  of a goodness-of-fit test of the theory  $\Theta$ , assuming measured values  $\mathbf{Q}$  of the observables:

$$p(\mathbf{Q}) = \int d^n \mathbf{Q}' \pi_{\hat{\mathbf{Q}}(\mathbf{Q})}(\mathbf{Q}') \theta(d^2(\mathbf{Q}') - d^2(\mathbf{Q})) \quad , \quad (8)$$

where  $\theta$  is the Heavyside step-function. Eq. 8 defines  $p(\mathbf{Q})$  as the probability that a random observable-vector  $\mathbf{Q}'$  distributed according to  $\pi_{\hat{\mathbf{Q}}(\mathbf{Q})}$  disagrees more with the theory  $\Theta$  than the observable-vector  $\mathbf{Q}$  which was actually measured. By using  $\pi_{\hat{\mathbf{Q}}(\mathbf{Q})}$  as PDF for  $\mathbf{Q}'$  we assume that the true observables are the best-fit observables  $\hat{\mathbf{Q}}(\mathbf{Q})$ , i.e. that the theory is realised with parameters from the best-fit domain  $\hat{X}(\mathbf{Q})$ . This is our *null hypothesis*.

To compare the performance of two different models one might be tempted to simply perform two goodness-of-fit tests and compare the resulting  $p$ -values. However, it is easy to construct an example where this strategy fails: Assume that we have 100 observables in our fit and model A (at it's best-fit point) describes the first 99 very well, but is in strong disagreement with the last observable. Model B, on the other hand, describes the first 99 observables as well as model A and also nicely agrees with the last observable. Then model B clearly performs better than model A, but the  $p$ -values of the two models would be essentially the same due to the large number of contributions to the log-likelihood function.

A broadly accepted strategy for comparing the performance of two models are likelihood ratio tests. This type of hypothesis test is used when one is not interested in the agreement of the theory as a whole with certain data, but rather in the agreement of some *restricted version* of a theory, assuming that the full theory is correct and describes the data. To this end, let us split up the parameter vector  $\mathbf{x}$  of the theory into a *relevant part*  $\mathbf{a}$  and an *irrelevant part*  $\mathbf{y}$ , i.e. write  $x = (\mathbf{a}, \mathbf{y}) \in A \times Y \equiv X$ , with the understanding that we are interested in the values of the parameters  $\mathbf{a}$  but not in the values of the parameters  $\mathbf{y}$ . For a fixed value of  $\mathbf{a}$  we now define the restricted versions  $\Theta_{\mathbf{a}}$ ,  $M_{\mathbf{a}}$ ,  $d_{\mathbf{a}}$  and  $\hat{\mathbf{Q}}_{\mathbf{a}}$  of the theory function, theory manifold, squared distance function, and best-fit observables as follows:

$$\begin{aligned} \Theta_{\mathbf{a}}(\mathbf{y}) &= \Theta(\mathbf{a}, \mathbf{y}) \quad , \quad M_{\mathbf{a}} = \{\Theta_{\mathbf{a}}(\mathbf{y}) | \mathbf{y} \in Y\} \subset M \quad , \\ d_{\mathbf{a}}^2(\mathbf{Q}) &= \min_{\mathbf{Q}' \in M_{\mathbf{a}}} L(\mathbf{Q}', \mathbf{Q}) \quad , \quad L(\hat{\mathbf{Q}}_{\mathbf{a}}(\mathbf{Q}), \mathbf{Q}) = d_{\mathbf{a}}^2(\mathbf{Q}) \quad . \end{aligned} \quad (9)$$

When comparing the restricted theory with the full theory one uses the *difference* of  $\chi^2$ -values of the restricted and the full theory as a test-statistic. To that end we denote the difference of squared distance functions as

$$\Delta_{\mathbf{a}}^2(\mathbf{Q}) = d_{\mathbf{a}}^2(\mathbf{Q}) - d^2(\mathbf{Q}) \quad . \quad (10)$$

Note that  $\Delta_{\mathbf{a}}^2$  is just a difference of minimal chi-square values for the measurement  $\mathbf{Q}$ , since  $d_{\mathbf{a}}^2(\mathbf{Q})$  and  $d^2(\mathbf{Q})$  are the minimal chi-square values of the restricted and the full theory, respectively. With these notations, the  $p$ -value  $p_{\mathbf{a}}(\mathbf{Q})$  for fixed relevant parameters  $\mathbf{a}$  and assuming measured values  $\mathbf{Q}$  of the observables is defined as

$$p_{\mathbf{a}}(\mathbf{Q}) = \int d^n \mathbf{Q}' \pi_{\hat{\mathbf{Q}}_{\mathbf{a}}(\mathbf{Q})}(\mathbf{Q}') \theta(\Delta_{\mathbf{a}}^2(\mathbf{Q}') - \Delta_{\mathbf{a}}^2(\mathbf{Q})) \quad . \quad (11)$$

It is the probability that some random set of measurements  $\mathbf{Q}'$  distributed according to  $\pi_{\hat{\mathbf{Q}}_{\mathbf{a}}(\mathbf{Q})}$  leads to a larger increase in disagreement than the actually measured point  $\mathbf{Q}$  when going from the full to the restricted theory. By using  $\pi_{\hat{\mathbf{Q}}_{\mathbf{a}}(\mathbf{Q})}$  as PDF for  $\mathbf{Q}'$  we assume that the true observables are the best-fit observables of the restricted theory  $\Theta_{\mathbf{a}}$  (our null hypothesis). Note that the  $p$ -value  $p_{\mathbf{a}}$  is equal to 1 if the parameters  $\mathbf{a}$  are equal to their best-fit values for the full theory. In this sense, likelihood ratio tests measure the performance of the restricted model in relation to the full model: if the restricted model describes the data as well as the full model, the likelihood ratio test gives a  $p$ -value of 1.

In the situation discussed above one also speaks of *nested models*: the restricted theory can be obtained from the full theory by fixing some of its parameters. Geometrically this means that the manifold  $M_{\mathbf{a}}$  of the restricted theory is a subset of the full theory's manifold  $M$ . Generalisations of likelihood ratio tests for non-nested models have been discussed, for example, in [7; 8]. From the geometric point of view, the generalisation is rather straightforward, since there is a canonical way of combining two theories into one.

Let  $\Theta_1$  and  $\Theta_2$  be two theories with unrelated parameter spaces of (possibly) different dimension and let  $M_1$  and  $M_2$  be the corresponding theory manifolds. The manifold of the combined theory is then simply the union of the two theory manifolds:

$$M = M_1 \cup M_2 \quad \Rightarrow \quad M_1, M_2 \subset M \quad . \quad (12)$$

The definition of the  $p$ -value in likelihood ratio tests of nested models only depends on the manifolds of the models, but not on their parametrisation. By using either  $M_1$  or  $M_2$  as the restricted manifold ( $M_{\mathbf{a}}$ ) and  $M$  as the manifold of the full theory, the definition (11) can be used without modification. The squared distance function of the full theory is given by

$$d^2(\mathbf{Q}) = \min\{d_1^2(\mathbf{Q}), d_2^2(\mathbf{Q})\} \quad , \quad (13)$$

where  $d_1^2$  and  $d_2^2$  are the squared distance functions of the theories  $\Theta_1$  and  $\Theta_2$ , respectively. The differences  $\Delta_1^2$  and  $\Delta_2^2$  of squared distance functions for restrictions to the theories  $\Theta_1$  or  $\Theta_2$ , respectively, are

$$\Delta_1^2(\mathbf{Q}) = \max\{d_1^2(\mathbf{Q}) - d_2^2(\mathbf{Q}), 0\} \quad , \quad \Delta_2^2(\mathbf{Q}) = \max\{d_2^2(\mathbf{Q}) - d_1^2(\mathbf{Q}), 0\} \quad . \quad (14)$$

For the theory  $\Theta_i$  with the smaller  $\chi^2$  value the function  $\Delta_i^2$  is 0 and the  $p$ -value is therefore always 1. This is analogous to the case of nested models, where the  $p$ -value at the best-fit point of the full theory is always 1.

### 3 Analytical Formulae for $p$ -values

Now that we have agreed on the definitions of  $p$ -values in goodness-of-fit and likelihood ratio tests, let us try to compute them. If all else fails the  $p$ -values defined in the last section can be computed with Monte Carlo methods. This has already been outlined in [5]. In practice, these Monte Carlo simulations can be very time consuming. However, under certain assumptions it is possible to derive simple (and well-known) analytical formulae for the  $p$ -values [3]. To understand the range of applicability of these formulae it is best to re-derive them. The derivations I present here use the geometric picture established in the last section. The derivations will also be instructive for the next section, where I discuss the possibilities for improving the efficiency of numerical computations of  $p$ -values.

Consider a set of observables  $\mathbf{O} = (O_i)_{i=1,\dots,n}$  distributed about their best-fit values  $\hat{O}_i$  according to an  $n$ -dimensional Gaussian distribution with a variance matrix  $V$ . In a first step we choose suitable linear combinations of the  $O_i$  to define a set of independent (i.e. uncorrelated) *normalised* observables  $Q_i$  which are distributed according to normal distributions (i.e. with expectation values of zero and standard deviations of one). To this end, let  $U$  be an orthogonal matrix that diagonalises  $V$ :

$$UVU^T = \text{diag}(\sigma_1^2, \dots, \sigma_n^2) \quad . \quad (15)$$

Then we define

$$Q_i = \frac{1}{\sigma_i} \sum_{j=1}^n U_{ij}(O_j - \hat{O}_j) \quad . \quad (16)$$

The PDF of the observables  $Q_i$  and the corresponding log-likelihood function are then given by

$$\pi_0(\mathbf{Q}) = \frac{1}{(2\pi)^{n/2}} e^{-\frac{1}{2}|\mathbf{Q}|^2} \quad \Rightarrow \quad L(0, \mathbf{Q}) = n \ln(2\pi) + |\mathbf{Q}|^2 \quad . \quad (17)$$

So we see that, up to a constant term, the log-likelihood function is just the squared euclidean length of the normalised observable-vectors  $\mathbf{Q}$ . The constant drops out in the definition of the squared distance function  $d^2$  and the  $\chi^2$  function.

As explained in the last section, the theory defines some manifold  $M$  in the space of normalised observables  $\mathbf{Q}$ . If  $M$  is sufficiently flat near the origin and extends sufficiently far away from it we can approximate  $M$  in the vicinity of 0 (the best-fit point) as a hyperplane. The flatness requirement will usually be valid if the errors  $\sigma_i$  of the original observables  $O_i$  are small, because in the transition to normalised observables the theory manifold is stretched by factors  $1/\sigma_i$ . The assumption of large extension is valid as long as the theory function is sufficiently regular in the vicinity of the best-fit domain  $\hat{X}(\mathbf{Q})$  and if  $\hat{X}(\mathbf{Q})$  is not close to any borders the parameter space  $X$  might have. In the presence of (small) systematic errors the extra parameters introduced to describe them have only small allowed ranges, so that all points in  $X$  lie close to the border. Thus the simple expressions for  $p$ -values derived in the following paragraphs should *not* be used in the presence of systematic errors. In addition to that, extreme values for the best-fit parameters or a singular behaviour of the theory function near the best-fit point also indicate a failure of these expressions.

Note that approximating the theory manifold as flat in the vicinity of  $\hat{\mathbf{Q}}$  is *not* the same as taking the theory function  $\Theta$  to be linear in its arguments. We just neglect the curvature of the manifold  $M$  in the vicinity of 0, but the map from  $X$  to the hyperplane  $M$  may still be non-linear.


If the theory manifold  $M$  is a hyperplane of dimension  $k \leq n$  through the origin, we may write any observable-vector  $\mathbf{Q}'$  as  $\mathbf{Q}' = \mathbf{Q}'_{\parallel} + \mathbf{Q}'_{\perp}$  with  $\mathbf{Q}'_{\parallel} \in M$  and  $\mathbf{Q}'_{\perp}$  perpendicular to it. The squared distance  $d^2(\mathbf{Q}')$  between  $M$  and  $\mathbf{Q}'$  is then simply the squared length of  $\mathbf{Q}'_{\perp}$ :

$$d^2(\mathbf{Q}') = |\mathbf{Q}'_{\perp}|^2 \quad . \quad (18)$$

For a two-dimensional sample space and a one-dimensional theory manifold, this is depicted in Fig 1. The  $\theta$  function in (8) vanishes in the region between the thin dashed lines. Eq. 8 now simplifies to

$$\begin{aligned} p(\mathbf{Q}) &= \frac{1}{(2\pi)^{n/2}} \int d^k \mathbf{Q}'_{\parallel} e^{-\frac{1}{2}|\mathbf{Q}'_{\parallel}|^2} \int d^{n-k} \mathbf{Q}'_{\perp} e^{-\frac{1}{2}|\mathbf{Q}'_{\perp}|^2} \theta(|\mathbf{Q}'_{\perp}|^2 - d^2(\mathbf{Q})) \\ &= 1 - P_{(n-k)/2}(\tfrac{1}{2}d^2(\mathbf{Q})) \quad , \end{aligned} \quad (19)$$





**Figure 1:** Orthogonal decomposition of the sample space in a goodness-of-fit test. The blue colour indicates the probability density for the toy observable vector  $\mathbf{Q}'$ . The  $\theta$  function in (8) vanishes in the region between the thin dashed lines.

where

$$P_a(x) = \frac{1}{\Gamma(a)} \int_0^x dt t^{a-1} e^{-t} \quad (20)$$

is the normalised lower incomplete Gamma function. Note that for a fixed experimental input  $\mathbf{Q}$  the quantity  $d^2(\mathbf{Q})$  is usually denoted as  $\chi_{\min}^2$  as it is the absolute minimum of the chi-square function (7).

A formula for  $p$ -values in likelihood ratio tests of *nested* models can be derived in a similar way. Keeping a subset  $\mathbf{a} = (a_i)_{i=1,\dots,r}$  of theory parameters fixed defines a  $(k-r)$ -dimensional sub-manifold  $M_{\mathbf{a}}$  of the theory manifold  $M$ . Now we shift the best-fit point  $\hat{\mathbf{Q}}_{\mathbf{a}}(\mathbf{Q})$  of the restricted theory to the origin. In the vicinity of the origin, we again approximate  $M$  as a hyperplane and  $M_{\mathbf{a}}$  as a linear subspace of  $M$ . We decompose the vector  $\mathbf{Q}'$  into three orthogonal components

$$\mathbf{Q}' = \mathbf{Q}'_1 + \mathbf{Q}'_2 + \mathbf{Q}'_3 \quad (21)$$

with  $\mathbf{Q}'_1 \in M_{\mathbf{a}}$  and  $\mathbf{Q}'_1 + \mathbf{Q}'_2 \in M$ . The function  $\Delta_{\mathbf{a}}^2(\mathbf{Q}')$  in (11) is then


$$\Delta_{\mathbf{a}}(\mathbf{Q}') = d_{\mathbf{a}}^2(\mathbf{Q}') - d^2(\mathbf{Q}') = |\mathbf{Q}'_2 + \mathbf{Q}'_3|^2 - |\mathbf{Q}'_3|^2 = |\mathbf{Q}'_2|^2 \quad (22)$$

For a three-dimensional sample space, a two-dimensional theory manifold  $M$  and a one-dimensional theory manifold  $M_{\mathbf{a}}$  this is depicted in Fig. 2. The  $\theta$  function in (11) vanishes in the region between the planes indicated by the thin dashed lines. Thus, Eq. 11 reads

$$\begin{aligned} p_{\mathbf{a}}(\mathbf{Q}) &= \frac{1}{(2\pi)^{n/2}} \int d^{k-r} \mathbf{Q}'_1 e^{-\frac{1}{2}|\mathbf{Q}'_1|^2} \int d^{n-k} \mathbf{Q}'_3 e^{-\frac{1}{2}|\mathbf{Q}'_3|^2} \int d^r \mathbf{Q}'_2 e^{-\frac{1}{2}|\mathbf{Q}'_2|^2} \theta(|\mathbf{Q}'_2|^2 - \Delta_{\mathbf{a}}^2(\mathbf{Q})) \\ &= 1 - P_{r/2}(\tfrac{1}{2}\Delta_{\mathbf{a}}^2(\mathbf{Q})) \quad (23) \end{aligned}$$

For a fixed experimental input  $\mathbf{Q}$ , the quantity  $\Delta_{\mathbf{a}}^2(\mathbf{Q})$  is usually denoted as  $\Delta\chi^2(\mathbf{a})$ , since it is the difference of minimal chi-square values in the restricted and the full theory.

We see that the derivations of the analytical formulae (19) and (23) rely on the following assumptions:



**Figure 2:** Orthogonal decomposition of the sample space for a likelihood ratio test of nested models. The blue colour indicates the probability density for the toy observable vector  $\mathbf{Q}'$ . The  $\theta$  function in (11) vanishes in the region between the planes indicated by the thin dashed lines.

1. The experimental errors on all observables are Gaussian (with or without correlations).
2. The theory manifold (in the space of normalised observables) can be approximated as a hyperplane. This approximation is valid if
  - (a) the curvature of the theory manifold (in the space of normalised observables) can be neglected and
  - (b) the best fit point(s) is (are) far away from any borders the theory manifold(s) might have.

Assumption 2(a) is valid if the errors of the observables are small enough. Assumption 2(b) fails if the mapping between the parameter space and the theory manifold is singular in the vicinity of the best-fit point or if any of the best-fit parameters are close to their upper or lower limit. In particular, assumption 2(b) is invalid in the case of (small) systematic errors, if these errors are treated within the *RFit* scheme.

In the case of non-nested models, as discussed at the end of the last section, the theory manifold can at best be approximated as a union of two hyperplanes with (possibly) different dimensionality. In this case Eq. 23 is clearly not applicable.

## 4 Numerical Calculation of $p$ -values

In the last section we have seen that there are many realistic scenarios where the simple analytic formulae (19) and (23) are not valid. In this case one has to resort to numerical

integration methods to calculate  $p$ -values. Based on the geometric view established in the last two sections I will now discuss some methods for improving the efficiency of these integrations. I will focus on the computation of  $p$ -values in likelihood ratio test, both for nested and non-nested models. The basic idea behind these methods is to optimise the numerical integration for the case where the errors on the observables are Gaussian and the manifolds of the two models (nested or non-nested) do indeed form hyperplanes in the space of normalised observables. The important difference is that these integration methods do not *rely* on the assumptions listed at the end of the last section. They only work best if they are valid.

Consider a set of observables  $O_i$  ( $i = 1, \dots, n$ ) whose distribution about their true values  $\hat{\mathbf{O}}$  is described by a log-likelihood function  $L(\hat{\mathbf{O}}, \mathbf{O})$ . By definition, the log-likelihood function, for fixed  $\hat{\mathbf{O}}$  and considered as a function of  $\mathbf{O}$ , has its absolute minimum at  $\mathbf{O} = \hat{\mathbf{O}}$ . We write the second-order Taylor-expansion of  $L$  about it's minimum as

$$L(\hat{\mathbf{O}}, \mathbf{O}) = L(\hat{\mathbf{O}}, \hat{\mathbf{O}}) + \frac{1}{2}(\mathbf{O} - \hat{\mathbf{O}})^\top V^{-1}(\mathbf{O} - \hat{\mathbf{O}}) + \mathcal{O}(|\mathbf{O} - \hat{\mathbf{O}}|^3) \quad . \quad (24)$$

In this approximation, the PDF  $\pi_{\hat{\mathbf{O}}}(\mathbf{O})$  is always a Gaussian distribution with variance matrix  $V$ . Thus, for an arbitrary log-likelihood function, we may define the *normalised observables*  $\mathbf{Q}$  analogous to Sec. 3, where  $V$  is now the Hessian matrix of the likelihood function at its minimum. The image of  $\hat{\mathbf{O}}$  in the space of normalised observables is 0 and the log-likelihood function in terms of the normalised observables is (up to an irrelevant constant)

$$L(0, \mathbf{Q}) = |\mathbf{Q}|^2 + \mathcal{O}(|\mathbf{Q}|^3) \quad . \quad (25)$$

Now assume that the observables  $\hat{\mathbf{O}}$  are a possible prediction of some theory  $\Theta$ . Let  $M$  be the manifold of the theory in the space of normalised observables. This means that  $M$  contains the origin. In the vicinity of the origin we may now approximate the theory manifold  $M$  by its tangent hyperplane  $H$  at the origin. Numerically, the hyperplane  $H$  can be constructed from the derivatives of the theory function  $\Theta$  with respect to all its parameters. Since  $H$  contains the origin, it is a subspace of  $\mathbb{R}^n$ . Similarly, we may consider some restriction  $\Theta_{\mathbf{a}}$  of the theory  $\Theta$ , whose image also contains the observable vector  $\hat{\mathbf{O}}$ . In the space of normalised observables this theory defines another subspace  $H_{\mathbf{a}} \subset H$ . All this is fairly similar to the case discussed in Sec. 3. There, however, the quadratic approximation (24) for the log-likelihood function was assumed to be exact and the theory manifolds  $M$  and  $M_{\mathbf{a}}$  were identical to the tangent hyperplanes  $H$  and  $H_{\mathbf{a}}$ .

Now recall the expression (11) for the  $p$ -value in a likelihood ratio test:

$$p_{\mathbf{a}}(\mathbf{Q}) = \int d^n \mathbf{Q}' \pi_0(\mathbf{Q}') \theta(\Delta_{\mathbf{a}}^2(\mathbf{Q}') - \Delta_{\mathbf{a}}^2(\mathbf{Q})) \quad , \quad (26)$$

where

$$\pi_0(\mathbf{Q}) = \frac{1}{(2\pi)^{n/2}} \exp[-\frac{1}{2}|\mathbf{Q}|^2 + \mathcal{O}(|\mathbf{Q}|^3)] \quad (27)$$

is the PDF in the space of normalised observables. To calculate the integral (26) numerically via Monte-Carlo simulation, one generates  $N$  sample points  $\mathbf{Q}'_i$ , which are randomly distributed according to some PDF  $\rho(\mathbf{Q}')$ . The integral is then estimated by

$$p_{\mathbf{a}}(\mathbf{Q}) \approx \frac{1}{N} \sum_{i=1}^N \frac{\pi_0(\mathbf{Q}'_i)}{\rho(\mathbf{Q}'_i)} \theta(\Delta_{\mathbf{a}}^2(\mathbf{Q}'_i) - \Delta_{\mathbf{a}}^2(\mathbf{Q})) \quad . \quad (28)$$

To reduce the statistical error of this estimate one has to choose the function  $\rho$  as similar as possible to the integrand, so that the terms in the sum are (ideally) all of the same size.

The first, obvious thing to do is to choose  $\rho$  proportional to  $e^{-\frac{1}{2}|\mathbf{Q}'|^2}$ , so that this factor cancels between  $\rho$  and  $\pi_0$ . This means that the sample points  $\mathbf{Q}'_i$  should follow an  $n$ -dimensional normal distribution. To further improve the estimate, we have to inspect the argument of the  $\theta$  function. Note that  $\Delta_{\mathbf{a}}^2(\mathbf{Q})$  does not depend on the integration variable, and  $\Delta_{\mathbf{a}}^2(\mathbf{Q}') = d_{\mathbf{a}}^2(\mathbf{Q}') - d^2(\mathbf{Q}')$ . Up to terms of order  $|\mathbf{Q}'|^3$  the squared distance functions can be calculated by projecting the vector  $\mathbf{Q}'$  on the orthogonal complements  $H_{\mathbf{a}}^{\perp}$  and  $H^{\perp}$  of the subspaces  $H_{\mathbf{a}}$  and  $H$ , respectively. Just as in the derivation of (23), we write  $\mathbf{Q}'$  as the sum of three orthogonal vectors  $\mathbf{Q}'_1$ ,  $\mathbf{Q}'_2$  and  $\mathbf{Q}'_3$  with  $\mathbf{Q}'_1 \in H_{\mathbf{a}}$ ,  $\mathbf{Q}'_1 + \mathbf{Q}'_2 \in H$  and  $\mathbf{Q}'_3 \in H^{\perp}$ . Then


$$\Delta_{\mathbf{a}}^2(\mathbf{Q}') = |\mathbf{Q}'_2 + \mathbf{Q}'_3|^2 - |\mathbf{Q}'_3|^2 = |\mathbf{Q}'_2|^2 \quad . \quad (29)$$

Thus, if the theory manifolds were identical to the hyperplanes  $H_{\mathbf{a}}$  and  $H$ , the integrand in (26) would be zero in the *inner region* defined by  $|\mathbf{Q}'_2|^2 < \Delta_{\mathbf{a}}^2(\mathbf{Q})$  and we should not waste time by throwing sample points in that region. Since the theory manifolds are not actually hyperplanes we should be a bit more careful and also put a few sample points in the inner region. There are many ways to do this. My suggestion is to define  $\rho$  as

$$\rho(\mathbf{Q}') = e^{-\frac{1}{2}|\mathbf{Q}'_1 + \mathbf{Q}'_3|^2} \begin{cases} a|\mathbf{Q}'_2|^{\alpha} & , \quad |\mathbf{Q}'_2|^2 < \Delta_{\mathbf{a}}^2(\mathbf{Q}) \\ be^{-\frac{1}{2}|\mathbf{Q}'_2|^2} & , \quad |\mathbf{Q}'_2|^2 \geq \Delta_{\mathbf{a}}^2(\mathbf{Q}) \end{cases} \quad . \quad (30)$$

The parameters and  $a, b, \alpha \geq 0$  may be tuned to improve the efficiency of the numerical integration (subject, of course, to the constraint that the PDF  $\rho$  is properly normalised). The problem of actually generating points according to this distribution will be addressed in Sec. 5.

Now consider the case of non-nested models  $\Theta_1$  and  $\Theta_2$ . Let us assume that  $\Theta_1$  has the smaller  $\chi^2$  value and we want to compute the  $p$ -value of  $\Theta_2$ , assuming that one of the two models is correct. To do this, we assume that  $\Theta_2$  is realised with its best-fit parameters. This is our null hypothesis. In the space of normalised observables  $\mathbf{Q}$  the best-fit point of  $\Theta_2$  is again the origin, while the best-fit point of  $\Theta_1$  is some other vector  $\hat{\mathbf{Q}}_1$ . We may again approximate the two theory manifolds by their tangent hyperplanes  $H_1$  and  $H_2$  at the respective best-fit points. Thus,  $H_2$  contains the origin and is a subspace of  $\mathbb{R}^n$ .  $H_1$ , on the other hand, contains  $\hat{\mathbf{Q}}_1$  but not necessarily the



**Figure 3:** Orthogonal decomposition of a three-dimensional sample space for non-nested models. The tangent hyperplane  $H_1$  of model  $\Theta_1$  is one-dimensional and the tangent hyperplane  $H_2$  of model  $\Theta_2$  is two-dimensional. The thin dashed line is the projection of  $H_1$  onto  $H_2$ . The blue colour indicates the probability density for the toy observable vector  $\mathbf{Q}'$ .

origin, so it is not a subspace. This makes an estimation of the support of the  $\theta$  function more difficult. Let  $H_{21} \subset H_2$  be the subspace obtained by shifting  $H_1$  by  $-\hat{\mathbf{Q}}_1$  (so that it contains the origin) and projecting it onto  $H_2$ . Furthermore, let  $H_{22}$  be the orthogonal complement of  $H_{21}$  in  $H_2$ . The projection of  $H_1$  onto  $H_{22}$  is then a single point  $\mathbf{C}$ , which can be obtained by projecting  $\hat{\mathbf{Q}}_1$  onto  $H_{22}$ . Any vector  $\mathbf{Q}'$  can now be written as the sum of three orthogonal component vectors  $\mathbf{Q}'_1$ ,  $\mathbf{Q}'_2$  and  $\mathbf{Q}'_3$  with  $\mathbf{Q}'_1 \in H_{21}$ ,  $\mathbf{Q}'_2 \in H_{22}$  and  $\mathbf{Q}'_3 \in H_2^\perp$ . The distance between  $\mathbf{Q}'$  and  $H_1$  is larger than  $|\mathbf{C} - \mathbf{Q}'_2|$  since the projection of any vector pointing from  $\mathbf{Q}'$  to  $H_1$  onto the subspace  $H_{22}$  is  $\mathbf{C} - \mathbf{Q}'_2$ . For the special case, where  $H_{21}$ ,  $H_{22}$  and  $H_2^\perp$  are all one-dimensional, these definitions are summarised in Fig. 3.

Now let us assume for a moment that the theory manifolds are identical to the hyperplanes  $H_1$  and  $H_2$ . For the squared distances  $d_1^2(\mathbf{Q}')$  and  $d_2^2(\mathbf{Q}')$  we then have

$$d_2^2(\mathbf{Q}') = |\mathbf{Q}'_3|^2 \quad , \quad d_1^2(\mathbf{Q}') \geq |\mathbf{C} - \mathbf{Q}'_2|^2 \quad . \quad (31)$$

Thus, the  $\theta$  function  $\theta(\Delta_2^2(\mathbf{Q}') - \Delta_2^2(\mathbf{Q}))$  vanishes for

$$|\mathbf{Q}'_3|^2 < \Delta_2^2(\mathbf{Q}) + |\mathbf{C} - \mathbf{Q}'_2|^2 \quad . \quad (32)$$

Note, however, that contrary to the case of nested models the  $\theta$  function may also vanish outside this region, even if the approximation of the theory manifolds by hyperplanes is exact. Nonetheless, a good strategy for improving the efficiency of the integration is to construct a PDF  $\rho(\mathbf{Q}')$  which “avoids” this region:

$$\rho(\mathbf{Q}') = e^{-\frac{1}{2}|\mathbf{Q}'_1 + \mathbf{Q}'_2|} \begin{cases} a|\mathbf{Q}'_3|^\alpha & , \quad |\mathbf{Q}'_3|^2 < \Delta_2^2(\mathbf{Q}) + |\mathbf{C} - \mathbf{Q}'_2|^2 \\ be^{-\frac{1}{2}|\mathbf{Q}'_3|^2} & , \quad |\mathbf{Q}'_3|^2 \geq \Delta_2^2(\mathbf{Q}) + |\mathbf{C} - \mathbf{Q}'_2|^2 \end{cases} \quad . \quad (33)$$

Note that the relation between  $a$ ,  $b$  and  $\alpha$ , which assures that  $\rho$  is properly normalised now also depends on  $|\mathbf{Q}'_2|$ . But there are still two free parameters which can be tuned to improve the efficiency of the numerical integration. To generate an ensemble of points distributed according to  $\rho$  one should first fix the components  $\mathbf{Q}'_1$  and  $\mathbf{Q}'_2$  by generating an equivalent number of normally distributed random variables. The component  $\mathbf{Q}'_3$  can then be generated with the steps described in Sec. 5.

Both, for the case of nested and non-nested models, the efficiency of the integration can be improved further by *adaptive* integration techniques, where the shape of the sampling density  $\rho$  is tuned *during* the actual integration. For the adaptation, the implementation in *myFitter* uses the OmniComp/Dvegas package [9] by Nikolas Kauer, which implements the VEGAS algorithm [10] and was developed in the context of [11; 12]. Thanks to OmniComp, parallelised integration is fully supported.

## 5 Introducing *myFitter*

The ideas for the numerical computation of  $p$ -values outlined in the last section have been implemented in a publicly available code called *myFitter*. The source code is available at Hepforge [13]. Detailed documentation will be included in the source distribution. Here I just want to provide a brief description of the application programming interface (API) and discuss some details of the implementation.

*myFitter* is a C++ class library and makes extensive use of inheritance and polymorphism to separate the tasks of fitting a model to experimental data and computing  $p$ -values from the tasks of implementing the observables (as functions of the model’s parameters) or the log-likelihood function (as a function of the observables). The main classes the user will have to deal with are:

**Model** This is the base class for all models implemented by the user. It essentially represents the theory function  $\Theta$  from earlier sections, i.e. the map from the model’s parameter space to the space of observables (sample space). The base class provides functionality for storing “current” values of parameters, observables and derivatives of observables with respect to the parameters, setting ranges in which parameters are allowed to float or fixing them (so that they do not float at all). It can also randomly sample the parameter space and build up a dictionary of parameter values and the corresponding observable values. This dictionary can be used to find good starting points for numerical minimisations of  $\chi^2$  functions. To implement their own model, the user has to subclass **Model** and overload the method `calc()` which computes the observables based on the current values of the parameters. They may also overload the method `calc_deriv()`, which calculates the derivatives of all observables with respect to all parameters. The default implementation uses simple numerical differentiation.

**LikelihoodComponent** This is the base class for objects that represent terms in the log-likelihood function  $L$  (see Sec. 2). Each likelihood component represents the contribution from one or more observables  $O_i$  to the log-likelihood function. To calculate the value of the log-likelihood function, the contributions of all **LikelihoodComponent** objects are added up. This is done by another class, **LikelihoodFunction**, which acts as a container for **LikelihoodComponent** objects. Derived classes of **LikelihoodComponent** must overload the method `calc( $\hat{\mathbf{O}}$ ,  $\mathbf{O}$ )`, which takes two vectors as arguments (the first being the “predicted” values of the observables and the second being the “measured” ones) and returns the contribution of the term to the log-likelihood function. Additionally, the methods `calc_deriv()` and `get_hessian()` must be implemented, which calculate the derivatives with respect to the  $\hat{O}_i$  and the Hessian matrix for the minimum at  $\hat{\mathbf{O}} = \mathbf{O}$ . Implementations for the most common likelihood components are also available. These classes are: **GaussianLC** (for single observables with a Gaussian and possibly systematic errors), **AsymmetricGaussianLC** (for single observables with asymmetric Gaussian error bars and possibly systematic errors) and **CorrelatedGaussianLC** (for several observables with Gaussian errors and a correlation matrix).

**Fitter** Objects of this type are responsible for fitting the parameters of models (represented by **Model** objects) to experimental data (represented by a **LikelihoodFunction** object) and for computing  $p$ -values by numerical integration. Each **Fitter** object contains a **LikelihoodFunction** object which is accessible through the `likelihood_function()` method and must be “filled” with **LikelihoodComponent** objects before any fits can be done. Once the likelihood function is initialised, fits can be performed with the `local_fit()` and `global_fit()` methods. As arguments, these methods take the **Model** object to be fitted and (optionally) a vector of central values for the observables. If no central values are given, the defaults from the `likelihood_function()` are used. The difference between these methods is that `local_fit()` uses the current values of the model parameters as starting point for the  $\chi^2$  minimisation, while `global_fit()` uses the dictionary created by a previous call to the model’s `scan()` method. The  $p$ -values for likelihood ratio tests of nested and non-nested models can be calculated with the methods `calc_nested_lrt_pvalue()` and `calc_lrt_pvalue()`, respectively. As arguments, these two methods take the models to be compared. Note that, for `calc_nested_lrt_pvalue()` to work, the second model must be a restricted version of the first, i.e. a copy of the first object with some additional parameters fixed. In addition to these methods, the **Fitter** class contains numerous options and flags that control the accuracy and various other aspects of the minimisation and integration routines. These options will be described in the package documentation. Most notably, the  $p$ -value integrations can be *parallelised without additional programming efforts* by the user.



To minimise the  $\chi^2$  function, *myFitter* currently uses the multidimensional minimisation algorithms from the GNU scientific library (GSL). The GSL provides several algorithms and the user can choose their favourite one with a call to `Fitter::minimizer()`. While these minimisation algorithms are sophisticated enough for most purposes, the possibilities for determining if a minimisation has converged are not. Currently, *myFitter* just aborts a minimisation when the length of the gradient is smaller than a number provided by the user. This is acceptable if the user also provides a good estimate of the width of the  $\chi^2$  minimum in the direction of each parameter. (These can be set with the `Model::scale()` method.) Other codes such as Minuit [14] use more sophisticated convergence tests and should be interfaced with *myFitter* in future releases.<sup>1</sup>

The problem of minimising a function of bounded parameters (i.e. of parameters that have an upper or lower limit) is solved in the usual way by smoothly and invertably mapping the real axis  $\mathbb{R}$  to the allowed range of the parameter. Internally, *myFitter* does this with the function

$$f : (-\infty, \infty) \rightarrow (0, \infty), \quad x \mapsto f(x) = \frac{1}{2}(x + \sqrt{x^2 + 1}) \quad . \quad (34)$$

However, since this way the region close to a parameter's upper or lower bound is stretched out to infinity, minimisation algorithms will often struggle to converge on a minimum that is close to or even on a parameter's upper or lower bound. Normally, some intervention by the user is required in this case. But when this happens during a  $p$ -value integration *myFitter* can not stop each time and ask the user what to do. It therefore uses a simple “snapping” rule: Whenever, during a  $\chi^2$  minimisation, a parameter gets too close to its upper or lower limit (the exact distance can be configured by the user), *myFitter* fixes that parameter at its limit and restarts the minimisation with one free parameter less.

Finally, let me explain the implementation of the integrand mappings from Sec. 4 in a bit more detail. To this end, note that the PDFs from Eq. 30 and 33 both decompose into three factors  $\rho_1(\mathbf{Q}'_1)\rho_2(\mathbf{Q}'_2)\rho_3(\mathbf{Q}'_3)$  where two factors are always proportional to an  $n_j$ -dimensional normal distribution ( $n_j$  being the dimension of the subspace that  $\mathbf{Q}'_j$  lives in). Generating vectors distributed according to an  $n_j$ -dimensional normal distribution is simple, because each component of the vector is then distributed according to a one-dimensional normal distribution.

The remaining problem is to generate an  $n$ -dimensional random vector  $\mathbf{Q}$  distributed according to the PDF

$$\rho(\mathbf{Q}) = \begin{cases} a|\mathbf{Q}|^\alpha & , \quad |\mathbf{Q}|^2 < \Delta^2 \\ be^{-\frac{1}{2}|\mathbf{Q}|^2} & , \quad |\mathbf{Q}|^2 \geq \Delta^2 \end{cases} \quad (35)$$

---

<sup>1</sup>In fact, I initially used Minuit2 [15] (the C++ version of Minuit) for minimisation, but the implementation proved unstable and would, in some cases, inexplicably start calling the  $\chi^2$  function with NaNs as parameter values.



for some  $\Delta > 0$ . We first note that the PDF  $\rho$  is rotationally invariant. The length  $R$  of the vector  $\mathbf{Q}$  is then distributed according to a PDF

$$\tilde{\rho}(R) = \frac{2\pi^{n/2}}{\Gamma(n/2)} R^{n-1} \begin{cases} aR^\alpha & , \quad R^2 < \Delta^2 \\ be^{-\frac{1}{2}R^2} & , \quad R^2 \geq \Delta^2 \end{cases} . \quad (36)$$

We may write this as

$$\tilde{\rho}(R) = f\tilde{\rho}_<(R) + (1-f)\tilde{\rho}_>(R) \quad (37)$$

where  $f \in [0, 1]$  is a free parameter and

$$\tilde{\rho}_<(R) = \frac{n+\alpha}{\Delta^{n+\alpha}} \theta(R) \theta(\Delta - R) \quad , \quad \tilde{\rho}_>(R) = \frac{R^{n-1} e^{-\frac{1}{2}R^2} \theta(R - \Delta)}{2^{(n-2)/2} \Gamma(n/2) (1 - P_{n/2}(\frac{1}{2}\Delta^2))} \quad (38)$$

are PDFs normalised to 1. Here,  $P_{n/2}$  is the normalised lower incomplete Gamma function (20). Since  $\tilde{\rho}_<$  and  $\tilde{\rho}_>$  are normalised, the parameter  $f$  is just the fraction of sample points that will be put in the inner region with  $R < \Delta$ . For a given  $f$ , the corresponding values of  $a$  and  $b$  are

$$a = \frac{f\Gamma(n/2)(n+\alpha)}{2\pi^{n/2}\Delta^{n+\alpha}} \quad , \quad b = \frac{1-f}{(2\pi)^{n/2}(1 - P_{n/2}(\frac{1}{2}\Delta^2))} . \quad (39)$$

By integrating  $\tilde{\rho}_<$ ,  $\tilde{\rho}_>$  and  $\tilde{\rho}$  from 0 to  $R$  we obtain the *cumulative distribution functions* (CDFs)

$$\begin{aligned} \text{CDF}_{\tilde{\rho}_<}(R) &= \frac{R^{n+\alpha}}{\Delta^{n+\alpha}} \quad , \quad \text{CDF}_{\tilde{\rho}_>}(R) = \frac{P_{n/2}(\frac{1}{2}R^2) - P_{n/2}(\frac{1}{2}\Delta^2)}{1 - P_{n/2}(\frac{1}{2}\Delta^2)} \\ \Rightarrow \quad \text{CDF}_{\tilde{\rho}}(R) &= \begin{cases} f \text{CDF}_{\tilde{\rho}_<}(R) & , \quad R < \Delta \\ f + (1-f) \text{CDF}_{\tilde{\rho}_>}(R) & , \quad R \geq \Delta \end{cases} . \end{aligned} \quad (40)$$

To generate random variables  $R$  distributed according to  $\tilde{\rho}$  we need the inverse of  $\text{CDF}_{\tilde{\rho}}$ :

$$\begin{aligned} \text{CDF}_{\tilde{\rho}_<}^{-1}(x) &= \Delta x^{1/(n+\alpha)} \quad , \quad \text{CDF}_{\tilde{\rho}_>}^{-1}(x) = \sqrt{2P_{n/2}^{-1}(x + (1-x)P_{n/2}(\frac{1}{2}\Delta^2))} \\ \Rightarrow \quad \text{CDF}_{\tilde{\rho}}^{-1}(R) &= \begin{cases} \text{CDF}_{\tilde{\rho}_<}^{-1}(\frac{x}{f}) & , \quad x < f \\ \text{CDF}_{\tilde{\rho}_>}^{-1}(\frac{x-f}{1-f}) & , \quad x \geq f \end{cases} , \end{aligned} \quad (41)$$

where  $P_{n/2}^{-1}$  is the inverse of the normalised lower incomplete Gamma function. This function is implemented in special functions libraries such as C++ Boost::math. Random vectors  $\mathbf{Q}$  distributed according to  $\rho$  may now be generated in the following way: First generate a vector  $\mathbf{Q}'$  according to an  $n$ -dimensional normal distribution. Then pick a uniformly distributed random variable  $x \in [0, 1]$  and set  $R = \text{CDF}_{\tilde{\rho}}^{-1}(x)$ . The variable  $R$  is then distributed according to  $\tilde{\rho}$ . The vector  $\mathbf{Q}$  with the correct random distribution is  $\mathbf{Q} = (R/|\mathbf{Q}'|)\mathbf{Q}'$ .

In this form, the integral is well-suited for *adaptive* integration by standard implementations of the VEGAS algorithm. Starting with a uniform distribution, we let VEGAS adapt the distribution of the variable  $x$  to improve the efficiency of the integration. Currently, *myFitter* does not use adaptation for any other integration variables, but this may change in the future. As mentioned earlier, the VEGAS implementation used by *myFitter* is the the OmniComp/Dvegas package [9] by Nikolas Kauer and supports parallelised integration.

## 6 Conclusions

To fully exploit the constraints that the LHC (and other) experiments impose on models of new physics, reliable methods for the computation of  $p$ -values in (frequentist) statistical analyses are needed. There are many realistic situations in which the well-known analytical expressions relating  $\chi^2$  or  $\Delta\chi^2$  values and  $p$ -values are not (necessarily) valid. These include the comparison of non-nested models and fits where systematic errors are treated within the *RFit* scheme. In these situations one has to resort to numerical methods.

Numerical computations of  $p$ -values are essentially a Monte-Carlo integration with an integrand function that is (usually) similar to a multidimensional Gaussian distribution times a  $\theta$  function which is zero in some *inner region* around the maximum of the Gaussian distribution. For small  $p$ -values, the inner region is very large and the Monte-Carlo integration becomes inefficient unless the sampling density  $\rho$  is chosen in a way that avoids the inner region. In this paper I have shown how to construct suitable sampling densities  $\rho$  in a very generic setting. These concepts were implemented in a C++ framework for maximum likelihood fits called *myFitter*, which is available for download at Hepforge [13].

The methods presented here and their implementation in *myFitter* have recently been applied in a global analysis [6] of the Standard Model with a fourth generation of fermions, where the analytical computation of  $p$ -values is impossible due to the non-decoupling nature of the fourth generation fermions.

## Acknowledgements

I would like to thank Jérôme Charles for fruitful discussions about likelihood ratio tests for non-nested models. I also thank Otto Eberhardt for checking fit results with CKMfitter and Ulrich Nierste for thorough proof reading.

## References

- [1] G. Casella and R. L. Berger, *Statistical Inference*, 2nd ed. (Duxbury Press, 2001).
- [2] K. Nakamura *et al.* (Particle Data Group), **J.Phys.G** **G37**, 075021 (2010).
- [3] S. S. Wilks, *Ann. Math. Statist.* **9**, 60 (1938).
- [4] O. Eberhardt, G. Herbert, H. Lacker, A. Lenz, A. Menzel, *et al.*(2012), [arXiv:1204.3872 \[hep-ph\]](#).
- [5] A. Höcker, H. Lacker, S. Laplace, and F. Le Diberder, **Eur.Phys.J. C****21**, 225 (2001), [arXiv:hep-ph/0104062 \[hep-ph\]](#).
- [6] O. Eberhardt, A. Lenz, A. Menzel, U. Nierste, and M. Wiebusch(2012), [arXiv:1207.0438 \[hep-ph\]](#).
- [7] Q. H. Vuong, *Econometrica* **57**, 307 (1989).
- [8] F. Lewis, A. Butler, and L. Gilbert, *Methods in Ecology and Evolution* **2**, 155 (2011).
- [9] <http://dvegas.hepforge.org>.
- [10] G. P. Lepage, **Journal of Computational Physics** **27**, 192 (1978), ISSN 0021-9991.
- [11] N. Kauer and D. Zeppenfeld, **Phys.Rev. D****65**, 014021 (2002), [arXiv:hep-ph/0107181 \[hep-ph\]](#).
- [12] N. Kauer, **Phys.Rev. D****67**, 054013 (2003), [arXiv:hep-ph/0212091 \[hep-ph\]](#).
- [13] <http://myfitter.hepforge.org>.
- [14] F. James and M. Roos, **Comput.Phys.Commun.** **10**, 343 (1975).
- [15] <http://seal.web.cern.ch/seal/snapshot/work-packages/mathlibs/minuit/>.